Server Encyclopedia

intel®

The Server Encyclopedia

Buying or selling servers requires in-depth comprehension of complex technologies. Understanding these technologies in the rapidly changing server industry requires continuous training and education. Intel designed this simple reference guide to help keep you and your customers informed regarding concepts and developments behind servers and networking. This guide will introduce you to topics such as reliability, availability, scalability, and benchmarks, giving you an inside look into server terms and concepts. This booklet has been designed to give you a brief overview on each of these subjects. Visit the Server Encyclopedia on the web to view or download more comprehensive white papers on the topics introduced in this booklet.

http://channel.intel.com/business/ibp/servers/encyc.htm

For server building blocks products from Intel:

http://developer.intel.com/design/servers

Table of Contents

Analysts in the Information Technology (IT) industry provide services to both consumers and vendors of IT. Consumers pay analysts for the information they provide in helping them invest in the right technology. Vendors pay analysts for their product and industry knowledge, often provided in the form of publications, training and consulting. *(This list is provided as a convenience only. It is not intended to be all inclusive.)*

Gartner Group

Gartner Group is one of the largest and best-known firms with over 400 analysts in 75 offices worldwide. They provide research and analysis of significant industry developments and trends and package their analysis into subscription-based products called personal advisory services. They also provide training and sponsor conferences.

IDC

International Data Corporation (IDC) is well known for its research and industry data. IDC has over 300 research analysts in forty countries who analyze vendors, assess technology, perform user surveys and interpret market trends.

Aberdeen Group

Aberdeen provides market research, analysis and consulting services, including helping clients develop information architectures and roadmaps, and helping them evaluate IT suppliers. Its publications include reports on industry events and trends, appraisals of companies and their products, and buying guides.

Patricia Seybold

The Seybold Group provides research, consulting and information services to both end-users and vendors. Its services include on-line interactive electronic service, technology-specific publications, case studies and IT architecture roadmaps.

D.H. Brown

D.H. Brown differentiates itself by concentrating on product research in terms of process features and functions. They evaluate how well product offerings and underlying technologies satisfy end-user needs. Their services include Product and Technology Evaluation Programs (PTEP), training and best practices implementations.

TPC

Run by the Transaction Processing Council (TPC), TPC benchmarks are published with a three-page executive summary and a full disclosure report including details on system configuration, system cost and benchmark methodology. The current TPC benchmarks are:

- TPC-C — measures complex On Line Transaction Processing (OLTP). This benchmark yields a number of transactions per minute (TPM). It also yields a price/performance metric derived by taking the price of the entire system and dividing by the performance, which equals a price per tpmC.
- TPC-D — measures Decision Support System (DSS) and measures applications that require complex, long running queries against large complex databases.

ServerBench

ServerBench is a Ziff-Davis benchmark which measures the performance of a server in a client/server environment. It measures data and applications running on a server, in addition to PC clients running Windows* 95 or Windows NT*.

NetBench

NetBench is another Ziff-Davis benchmark that measures how well a file server handles file I/O requests from as many as four different client operating systems (DOS, 32-bit Windows, 16-bit Windows* and/or Mac* OS). NetBench totals all the client throughput (based upon network file operations) to produce the overall server throughput.

NetBench differs from ServerBench in that NetBench measures the performance of file servers, while ServerBench measures the performance of application servers.

SAP

The SAP benchmarks measure server products that run the SAP R/3 software application environment. The tests measure performance of a specific client/server system configuration combining server(s), relational database and SAP R/3 software. The tests are designed to perform the most typical transactions of the business model.

Tower

A tower is a computer cabinet that is taller than it is wide (also called a deskside computer).

An advantage of the tower chassis over desktop computers is the small "footprint" (the amount of floor space it uses). A tower is generally 18 to 27 inches tall. A mini-tower is approximately 14 inches tall, and a mid-sized tower is generally around 16 inches tall.

Pedestal

The Pedestal is a wider and often taller version of the tower. Pedestals offer more flexible configuration options and scalability than towers and often provide hot swap capabilities and disk arrays.

Rack Mount

A rack mounted system allows nodes or chassis to be mounted onto a metal frame where they can slide on rails.

A typical rack is 77 inches tall, 24 inches wide, and 40 inches deep. A rack is measured in units of vertical measurement, represented by the letter "U." One "U" = 1.75 inches or 4.445 centimeters. A 77-inch vertical rack is 40U.

The rack system may consist of monitors, disk drives, uninterruptible power supplies (UPS), network components and server nodes.

An extension of the rack chassis concept breaks up the server node into separate pieces that communicate through a server area network. There may be separate processing nodes, memory nodes, expansion bus nodes and disk arrays.

Server Clustering

A cluster is two or more computers (nodes) in a group that provide higher availability and scalability than would be possible if the computers worked separately.

Availability

When any node fails, its resources can "fail over" to one or more other nodes in the cluster, thus improving the availability of all the nodes. All nodes are connected by one or more "heartbeat" paths between systems. If the heartbeat dies, the failover process starts.

Scalability

Clusters increase scalability by allowing the addition of processing power or disk capacity without interruption of service. Additional scalability will be seen in the future as standards emerge for developing "cluster aware" applications that can automatically spread their workload across multiple nodes in the cluster.

Cluster Configurations

Shared storage — A common arrangement. Each node uses local disks to store operating system swap space and system files, the application data is stored on the shared disks. Each node can read the data written by other nodes.

Server mirroring — Less expensive than shared storage. When data is written to the primary server, it is also written to a secondary server. Data integrity is maintained by locking the secondary server data. Some server mirroring products can also switch workload from primary to secondary server.

Shared nothing — Data is not shared. In a failure situation, the shared-nothing cluster has software that can transfer ownership of a disk from one node to another.

N-way configuration — All nodes are normally active with its own set of users and workload. A failed node's resources can failover to other nodes, which can cause some degradation of the remaining servers' performance as they take on additional workload.

N+1 configuration — Cluster includes a hot standby node that waits in idle mode for a primary system failure. A node failure avoids performance degradation in the other nodes, but the cost of overhead is higher since the standby is not normally an active service provider.

During Client/Server Application Development, system planners must focus on three basic parameters: quality, cost and time. These three parameters form a fixed equation where a change in one parameter *must* be offset by a change in one or both of the remaining parameters. For example, if time is reduced, then quality must go down or cost must increase. A successful project is one that is completed at the specified level of quality, on or before the deadline and within budget.

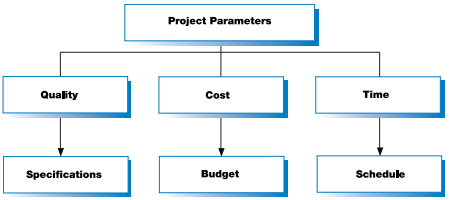

Layered Architectures

One approach to designing client/server systems is to focus on clearly defined layers of the application architecture. Client/server system layers can be designed independently by specialists provided that the connections or interfaces between the layers are carefully planned. Each layer must remain independent of the other layers with clearly defined domains and protocols for interacting with one another.

- Two-layer Architecture
  1. *Application layer* — handles presentation, application and transaction logic. Presentation logic is Graphical User Interface (GUI) with which the user interacts with the application. The application logic handles the business rules and policies. The transaction logic is the code that groups database updates into transactions and helps to ensure that all updates within a transaction are made consistently.

2. *Database layer* — This layer consists of the underlying database engine that supports the application. It is responsible for maintaining the integrity of the database. Some or all of the transaction logic may be implemented in this layer.

- Three-layer Architecture
  1. *User Interface Layer* — design of the Graphical User Interface (GUI) logic.
  2. *Business Rule Layer* — contains the rules and policy logic unique to the application or organization.
  3. *Database Layer* — Maintains the integrity of the database. It is the underlying database model that supports the application.

Terminology

Prototyping — the creation of the shell of an application, giving the illusion of a complete or nearly complete program.

Rapid Application Development — attempts to shorten the overall development cycle through the use of prototyping and overlapping steps in the development process. Analysis, design, implementation and testing are all done simultaneously in incremental steps.

Joint Application Development — an approach in which all parties which may be involved at some point in the development of the application are brought together to work through the various issues. Includes designers, a high level manager, end-users, and technical specialists.

The Data Warehouse is a method of making the accumulated data of a business available to decision-makers. It should show trends, help to forecast the future of the business and aid in setting corporate strategy.

In any business organization there are two types of data, Operational Data and Analytical Data.

Operational Data consists of information that is tactical. It is data that is needed to run the day-to-day operations of a business. Examples include invoicing, account balances and inventory control.

Analytical Data is strategic information. It helps to identify trends, recognize opportunities and pinpoint problems. Analytical data is what Data Warehouses are all about. Analytical data must be sheltered from the constant updating and modification of operational data. Analytical data is not accurate to the moment, but is accurate over a specific period of time that could be a week, month or a year.

Data Warehouse Characteristics

Regardless of the make, model, or vendor, all good Data Warehouses will have the following characteristics:

- Separate from the operational database
- Based on a comprehensive model of the organization's information resources
- Data delivered from multiple operational databases
- Automatically maintained

- Accessible from many query tools
- Read-only
- Can be replicated, no need to keep information in a central location
- Information is presented in consistent and compatible formats (e.g., revenue is presented and converted to one currency such as U.S. Dollars or German Marks)

Terminology

Extract Programs — Data Warehouses get their data from operational databases. Periodically an extract program is run which moves detailed transaction data to the data warehouse.

Datamart — A subset of a Data Warehouse for a single department or function. It is usually much smaller than a data warehouse.

Decision Support System (DSS) — Extract and report software programs used by management to help make strategic decisions. Usually, it looks at a specific type of data to provide reports and may or may not be used in conjunction with a Data Warehouse. Provides the ability to model by answering "what if" types of questions.

Executive Information System (EIS) — Provides a concise snapshot of how a business is doing on a daily basis.

Fibre Channel is a storage and network protocol useful for high-speed network and storage requirements. Although it sounds as if it is restricted to optical fiber, it is, in fact, a standard that can be implemented using either copper or optical cabling. Physically, Fibre Channel uses two pairs of wires, one pair for sending and the other pair for receiving, both of which can operate simultaneously (called Full duplex).

Fibre Channel's primary advantages are its speed and flexibility. The high speed capability and low latency of Fibre Channel makes it ideal for applications which require large data transfers such as systems dedicated to imaging, 3-D rendering, video production and large data mining applications. Manufacturers are also offering Fibre Channel switches (known as Fabrics) and adapters as a total solution to storage and networking.

Fibre Channel:

- Protocol independent
- Supports data rates of 12.5 to 100 megabytes per second
- May use optical fiber, coax or twisted pair cabling
- May support up to 127 connections on the same loop

Fibre Channel Services

Service Class 1 — A hard-wired or circuit switched dedicated connection between two nodes that cannot be broken, similar to a telephone connection. It is ideal for use wherever there is a need for dedicated links, such as one between server and storage.

Service Class 2 — A connectionless, frame switched service that guarantees delivery with confirmation. Similar to packet switched technologies like frame relay, switching is performed on the data frame rather than on a connection, and frames are sent over any available route. It is best for environments that involve a mixture of storage and network communications.

Service Class 3 — Similar to Class 2, Class 3 is a connectionless frame switched service. Designed for one-to-many connections, it contains no guarantee or confirmation mechanism. There are no retransmissions or waits for acknowledgment, thereby enhancing performance. It is ideal for storage, which generally encompasses its own confirmation protocol.

Service Class 4 — A connection-based service that offers guaranteed fractional bandwidth and guaranteed latency levels. Users are able to lock down specific paths through a Fibre Channel switch. Class 4 supports isochronous service that can carry real time video and video traffic as well as data.

The I²C* Management Bus

The I²C* Management Bus is a two wire serial bus used for connecting microcontrollers and their peripheral devices. The name I²C is an acronym for inter-integrated circuit bus which literally explains its function: to provide a communication link between integrated circuits.

What Does it Do?

Originally developed for audio and video equipment in the 1980s, the I²C bus is now used primarily for server management, which involves communicating the status of individual components. For instance, queries can be made of various components to find out the configuration of the system or to learn the functional status of components, such as power supplies and system fans.

Advantages

A primary advantage of the I²C bus is its simplicity and efficiency. The bus takes up very little space because the interface is directly on the component. This leads to reduced board space and pin count, thereby lowering interconnectivity costs. The length of the bus can be up to 25 feet and may support as many as 40 components at a maximum transfer rate of 100 kilobits per second.

Another advantage of the I²C bus is that it supports multimastering, wherein any device capable of transmitting and receiving may become the busmaster. A busmaster can control signal transfer and clock frequency. Naturally, there can be only one master at any point in time.

I²C Bus Example

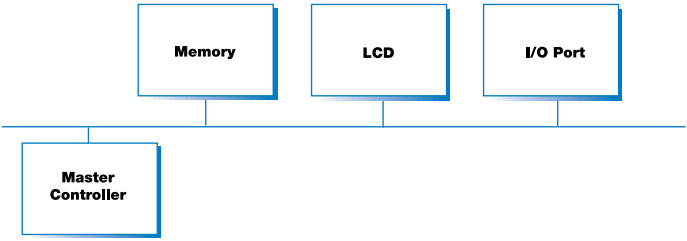

What is I_2O?

I_2O is a standard interface for intelligent I/O systems that works across diverse operating systems and software revisions. I_2O was developed in 1996 by Intel and other industry leaders to address the need for greater I/O throughput.

Why I_2O?

An interrupt occurs whenever a disk subsystem, a network interface card or other I/O device needs attention. In any given operation, an I/O device may interrupt the processor many times. While the processor chips used today are fast, they were not designed to handle interrupt duties. I_2O allows the processor to do what it does best, manage the applications and off-load the I/O functions by implementing intelligent I/O processing.

I_2O allows requests to come in from a device on PCI without going through the main processor. The I_2O host processor recognizes the requests and handles them locally. It also allows requests to queue up at the I_2O processor while the main processor is working on other tasks.

I_2O:

- Reduces the workload of the main processor resulting in improved system performance
- Increases I/O throughput
- Provides a standard interface to I/O devices
- Reduces the number of drivers needed for peripherals

How It Works

I_2O drivers are divided into two Modules: The OS Services Module (OSM) and the Hardware Device Module (HDM).

The OSM interfaces with the operating system and the HDM interfaces with the hardware device. The two modules exchange information through a two-layer communications system in which a message layer sets up a communications session, while a transport layer defines how information will be shared. The modules communicate without knowledge of underlying bus architectures or topologies. Messages take the form of meta-language, so that communications do not depend on host operating systems or bus configurations.

The Internet

The Internet consists of many, many networks interconnected very loosely. There is no single place where all the connections are registered. It originated as the U.S. military ARPANET and was helped along by the National Science Foundation (NSF). The Internet now spans the globe and reached over 19 million hosts in July of 1997! (see http://www.nw.com/zone)

The Internet is fundamentally a mechanism for sending and receiving information anywhere in the world. It uses the TCP/IP communications protocol. An intranet is an internet system confined within a company or organization by special software called a “firewall”.

Internet Services

The most commonly used Internet services are:

- E-mail
- FTP (file transfer protocol)
- Usenet news
- World Wide Web

World Wide Web

The Web is the part of the Internet that provides a way to get to all the information and services available on the Internet. These services are available through software called a web browser, which reads the standardized web language called HyperText Markup Language (HTML).

The Uniform Resource Locator (URL) is the address your browser uses to find information located on another computer and to retrieve the server's corresponding HTML pages. The following is an example of a URL: http://www.intel.com/home

A computer uses two types of memory:

- Main memory, often referred to as Dynamic Random Access Memory (DRAM), is inexpensive and relatively slow.
- Cache is a relatively small amount of memory which is costly and fast. Cache is referred to as Static Random Access Memory (SRAM). It speeds up processing by being placed near the CPU.

The CPU takes data from the disk and places it into main memory. Data that is being processed immediately is moved to first level cache.

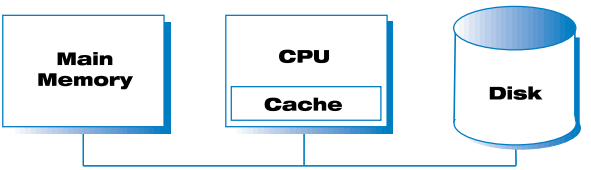

Memory Speeds, Feeds and Terms

Memory access speed is measured in nanoseconds. A nanosecond is one billionth of a second. Incredibly, memory access speeds can be as much as 1000 times faster than disk speeds.

DRAM

A typical access speed for main memory (DRAM) is about 60 nanoseconds (60ns).

SRAM

Static Random Access Memory (SRAM) is much faster than DRAM, but it is more expensive, which is why most systems have large amounts of DRAM and small amounts of cache (SRAM).

SIMMs

Single In-line Memory Modules come in different density sizes, and are placed on a DRAM memory board. SIMMs use a 64-bit data path.

DIMMs

Dual In-line Memory Modules utilize both sides of the memory board, providing up to double the memory of SIMMs using the same amount of board space. DIMMs use a 128-bit data path.

Interleaving

Interleaved memory is like parallel processing for memory. It allows consecutive addresses to reside in different blocks that can be handled simultaneously, which improves bandwidth and performance.

The two most common types of network architecture are Ethernet and Token Ring.

Ethernet*

The physical characteristics and the speed of the cabling distinguish the different types of Ethernet networking architecture.

Ethernet networks transmit only when the channel is quiet. This technique is called carrier-sense multiple access with collision detection (CSMA/CD).

Benefit:
Allows relatively fast access when fewer, longer data frames are transmitted.

Drawback:
Collision management and retransmission can degrade performance when short bursts of data are being transmitted.

There are three main types of Ethernet cabling: thinnet, twisted pair and fast Ethernet or Gigabit.

Thinnet:
Named from the thin coaxial cables used to connect the network, thinnet is an inexpensive and economical configuration, ideal for a small department or workgroup.

Twisted Pair:
The twisting cancels electrical noise from adjacent pairs and other devices such as motors and transformers.

Fast Ethernet and Gigabit Ethernet:

The latest and next generation of Ethernet (gigabit is just being implemented) provides a tremendous advantage in scalability because all Ethernet types use the same frame format, duplex operations and flow control methods.

Token Ring

Distinguished by its token passing communication method rather than by the physical characteristics of its cabling. The main types of token ring networks are ARCnet, FDDI and Circuit Switched Digital Services.

ARCnet:

(Attached Resource Computer Network) is easy to install and inexpensive. ARCnet was originally limited by its signaling speed of 2.5 Mbits/second, but the newer version ARCnet Plus provides 20 Mbits/second throughput.

FDDI:

(Fiber Distributed Data Interface) can transmit at 100 Mbits/second at a distance of up to two kilometers. FDDI provides communication services and acts as a traffic gathering system.

Circuit Switched Digital Services:

An alternative to high cost dedicated LAN-to-LAN, dial-up circuit vendors offer signaling rates from 56kbs to 1.544 Mbps.

A computer network is two or more computers connected together in order to share information. Networked computers perform three key functions: they share files, resources (printers, disks, etc.), and programs.

Topology

Topology refers to the way a network is wired or arranged. The three most widely used topologies are Bus, Star and Token Ring.

Bus

The simplest type of network that connects computers in a row along a single cable segment. An Ethernet network is based on a bus topology.

Advantages	Disadvantages
Failure of single computer doesn't affect the entire LAN	Cable break may affect large number of users
Easy to connect	Cable length is limited
Inexpensive	Difficult to isolate cable errors

Star

A star topology connects computers to a central hub with each computer having its own connection.

Advantages	Disadvantages
Easy to add workstations	Hub failure results in loss of communication for all computers attached to the hub
Simplifies network management	

Token Ring

Computers are connected in a continuous network loop in which a token is passed from one computer to the next. The token is a data frame (or packet) which is continuously passed around the ring.

Advantages	Disadvantages
Cable failure affects few users	Physical connections are more costly
Equal access for all users	
Easy to add workstations without reduced network performance	

The common network hardware components include hubs, media (or multistation) access units (MAUs), repeaters, bridges, routers, gateways, network interface cards (NICs) and cabling.

Hub

Hubs can be active or passive. Passive hubs just listen. Active hubs listen and regenerate the data in order to maintain signal strength.

Media (or Multistation) Access Unit

MAU's are used only in Token Ring configurations. They operate much like a hub and can also be either active or passive.

Repeater

Physically small and used to connect two segments of network cable, repeaters are the simplest and least expensive network hardware.

Bridge

A bridge is an intelligent repeater that listens to the network traffic and directs messages to workstations on either side of the bridge. [Similar Networks]

Router

A router is a very smart bridge. It performs the same tasks as repeaters and bridges, but can also determine the most efficient path to transmit the data.

Gateway

Gateways connect radically different networks together. Although slower than a bridge or a router, its own processor allows for translation of different protocol between the networks.

Network Interface Card

A NIC card or LAN adapter acts as an interface between the computer and the network cabling.

Cabling

Cable	Speed	Cost	Diameter	Distance
Unshielded Twisted Pair	Fast enough	Least expensive	Small	Short
Coaxial	Very fast	Inexpensive	Medium	Medium
Shielded Twisted Pair	Very fast	Expensive	Large	Short
Fiber-optic	Fastest	Most expensive	Small	Very Long

The PCI (Peripheral Component Interconnect) bus was jointly developed by Intel and other industry leaders in order to bring current and next-generation PCs to new levels of systems performance. The PCI bus takes peripherals off the I/O bus and connects them, together with the CPU and the memory subsystem, to a wider, faster pathway for data. PCI began as an advanced high-performance local bus that supports multiple peripherals and has grown into an I/O bus and expansion bus.

How PCI Works

PCI occupies an intermediate level between the CPU local bus (processor/memory/cache subsystems) and a standard expansion bus (ISA, EISA, MicroChannel). The PCI bus is isolated from the CPU local bus by a PCI bridge/controller. The CPU can write data to peripherals such as a hard drive and the PCI bridge/controller can immediately store the data in its buffer. This lets the CPU go on to its next operation rather than waiting for the transfer to complete. The buffer then feeds the data to the peripheral at the most efficient rate possible.

Advantages of PCI

- High Performance — greater than 100 MB sustained
- Low Cost
- Longevity
- Software and hardware compatibility

BUS Comparison Chart

	ISA	EISA	MicroChannel	VESA	PCI
Data Path Width	8/16	32.00	16/32/64	32	32/64
Data Bus Speed (MHz)	5.33/8.33	8.33	10.00	33/50	33.00
Data Transfer Rate (MB/Sec)	5.33/8.33	33.00	20/40/80/160	132	132/264
Data Rates (MB/Sec)	5.33/8.33	33.00	80 (RS/6000)	132.00	132.00
Number of Slots	0 - 8	0 - 8	0 - 8	0 - 2	0 - 4
Busmaster Supported	NO	YES	YES	YES	YES
Data/Address Parity	NO	NO	YES	NO	YES
Syn, Channel Checks	NO	NO	YES	NO	YES
Card ID/ Auto Config.	NO	YES	YES	YES	YES
Works with MC/ISA/EISA	N/A	N/A	N/A	YES	YES

RAID was designed to revolutionize the way computers manage and access mass storage of data by providing an inexpensive and redundant system of disks.

RAID accomplishes its goals of redundancy and fault tolerance by striping, parity checking and mirroring.

- *Striping* means that files are written a block at a time over multiple disks. The striping technique divides data across many drives and improves data transfer rates and total disk transaction times.
- *Parity checking* ensures that the data is valid by performing a redundancy check on all data following a transmission. With parity, one of the disks on a RAID system can fail and the other disks have the ability to rebuild the failed disk.
- *Mirroring* involves the creation of a duplicate copy of a disk. Mirroring is a technique where the data is written simultaneously to a pair of drives. The penalty paid is that two drives must be purchased to get the capacity of only one. Mirroring provides the same general protection as parity, but has more overhead.

RAID Levels

- RAID 0 — Data striping array
- RAID 1— Mirrored disk array

- RAID 2 — Parallel array, hamming code
- RAID 3 — Parallel array with parity
- RAID 4 — Independent actuators with a dedicated parity drive
- RAID 5 — Independent actuators with parity spread across all drives

Components

The major components in RAID are the Disk Array Controller (DAC) and a rank of five disks. The picture below shows an example of RAID 5. Data is striped across all five disks and the parity is used to recover a failed disk.

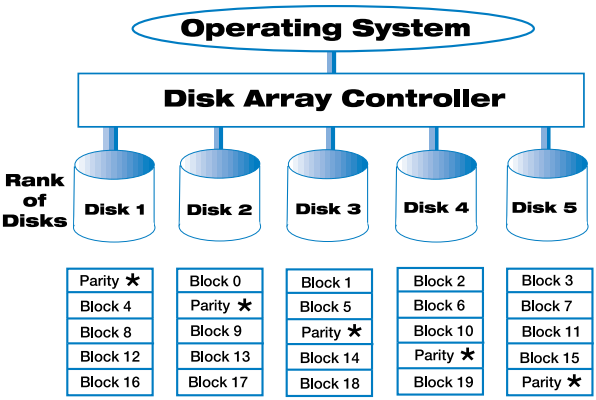

*Parity spread over multiple disks

RAID 5 Example

The relational database is any database that stores data in tables that can establish relationships with other tables based on common information.

Data Distribution Methods

Downloading — Often the simplest and best method of data distribution is to reproduce it by Downloading. Downloading is best used in situations where data is not considered mission critical on a daily basis.

Replication — Replication is the ability to duplicate key portions of a database in various locations, making sure that all copies of the data are simultaneously updated. Replication is best used where the users must have the most current information.

Fragmentation — Fragmentation splits up portions of a table among several servers. Tables may be split horizontally or vertically. Horizontal Fragmentation splits the table into two or more groups of rows and stores the rows on separate servers. Vertical Fragmentation splits a table into two or more columns and stores the columns on separate servers.

Data Integrity

Data Integrity refers to the validity and ability to trust the data. One of the ways in which Data Integrity is maintained is through the use of constants which establish values for specific fields, and Nulls which are fields that have no value or the value is unknown.

Referential Integrity

Referential Integrity checks the relationships of data contained in one table with data referenced in another table. Referential Integrity is enforced using Declarative Referential Integrity which builds constraints into the system when the database is initially designed and Triggers which are small programs that run whenever data is altered.

Declarative Referential Integrity can be thought of as proactive, whereas Triggers may be thought of as reactive.

Business Rules

Business Rules are unique to a particular organization and may reflect field validation and integrity. An example of a Business Rule is:

Orders are accepted with some credit cards but not others.

Concurrency and Locking

Concurrency allows multiple users to access a database at the same time by "locking" the data. A shared lock is a read-only lock that allows users to read the data but not update it. An exclusive lock prevents users from reading or updating the data until the lock is released.

Database Terminology

SQL — SQL stands for Structured Query Language and is the standard programming language used to develop and access relational databases.

Middleware — Middleware refers to Client/Server software that exists between the client and the server.

DBMS — Database Management System is the hardware and software on which the database is maintained.

Join — Retrieving and combining multiple tables.

Normalization — The process of eliminating redundant information in a database design.

OLTP — On Line Transaction Processing applications are the mission critical applications that handle day to day transactions, such as order processing, shipping, accounts payable, etc.

OLAP — On Line Analytical Processing is decision support software that allows the user to analyze information that has been summarized. OLAP tools are used to perform trend analysis.

System availability is influenced by four factors:

- Hardware (controllable)
- Software (controllable)
- People (more difficult to control)
- Environment (very difficult to control)

Even if it were possible to create hardware and software that never failed, human and environmental factors still create the need for redundancy and failover functions.

Reliability

Reliability is often measured as a function of the time between system failures using the term MTBF (Mean Time Between Failures). Computer and component builders usually quote MTBF in thousands of hours.

High Availability

Availability is measured as the percentage of time that a system is functioning and usable. Because 100% system availability is very costly to achieve, users should identify their mission critical applications and the cost of each hour of downtime for those applications. Then the user can determine the right combination of high availability features that will reach an acceptable level of availability.

Examples of high availability features are:

- RAID (Redundant Array of Independent Disks)
- Server management tools
- UPS (Uninterruptible Power Supplies)
- Powerfail recovery
- Redundant components
- Hot pluggable disks, fans and power supplies
- System or application failover

Fault Tolerance

The highest level of system availability is provided by fault tolerant (FT) systems which incorporate redundancy into virtually every aspect of the system. True fault tolerant systems tend to be very expensive and traditionally have not offered high performance.

Combining high availability features with microprocessor-based systems today allows users to reach higher and higher levels of availability at a more affordable cost, with no sacrifice in performance.

Scalability

A machine that is scalable has the ability to grow in size and speed. Some machines offer limited scalability by design, while some can grow to virtually any size needed.

Computers Scale by:

- Using multiple processors
- Adding memory
- Increasing disk storage

Symmetrical Multiprocessing

A Symmetric Multiprocessing (SMP) machine is a computer that uses multiple processors. These processors share memory and one copy of the operating system. SMP machines can scale by starting small with only two processors and adding more processors as business needs and applications grow.

Parallel Processing

Parallel processing takes SMP a step further by combining multiple SMP "nodes" that can work in parallel on a single application, usually a database that is fully "parallel-capable". Because each node has its own copy of the operating system, and the nodes communicate through a specialized interconnect, adding additional nodes does not tax a single operating system. Therefore, parallel processing can scale in a more linear fashion than SMP alone.

Upgrading vs. Scaling

What is the difference between an upgrade and scaling? An upgrade is the ability to replace a component with a faster or better component, whereas a machine that is scalable allows the user to add to or build on that component. For example, a PC normally has only one processor that a user might be able to upgrade with a simple replacement. The user cannot scale a processor because most PCs are not designed for two or more CPUs.

Factors That Limit Scalability

- Operating System may limit the number of processors it can support
- Applications may not be designed to take full advantage of additional system capabilities
- Memory may be shared among all processors
- Number of slots for cards and memory
- Limited physical expansion capabilities of chassis

The Small Computer System Interface, better known as SCSI (pronounced "scuzzy"), is an industry and ANSI standard which provides a host computer with an intelligent interface to peripherals. SCSI allows several disk drives, CD-ROM drives, scanners, printers and other external input/output devices to be accessed concurrently, simultaneously or in sequence. SCSI is considered an intelligent interface because it includes a processor that can handle the I/Os to and from a device.

The Major Components of SCSI:

- Host adapter card containing one controller and a terminator
- SCSI Bus (also called a SCSI Cable)
- SCSI-compatible peripheral devices (disks, CD-ROM, tapes, etc.)
- A terminator representing the end of the SCSI Bus

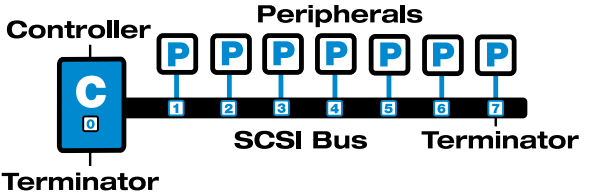

SCSI Specification Table

SCSI Type	Bus	Transfer Rate/sec.	Peripherals Supported
SCSI 1	8-bit	2.5 MB	8
Synchronous SCSI-2	8-bit	5.0 MB	8
Fast SCSI-2	8-bit	10 MB	8
Fast and Wide SCSI-2	16-bit	20 MB	16
Ultra SCSI	8-bit	20 MB	8
Ultra2 SCSI	8-bit	40 MB	8
Wide Ultra SCSI	16-bit	40 MB	16
Wide Ultra2 SCSI	16-bit	80 MB	16

Serviceability is the ease with which the hardware and software features of a system facilitate maintenance and repair functions. Serviceability has a direct effect on the total cost of ownership (TCO), which includes the purchase price, cost of administration and cost of maintenance.

There are four server characteristics that can affect serviceability, described below.

Redundancy

Servers should have spare components in any situation that is a single point of failure — that is, where one failure could cause the entire system to come down. Disks, fans and power supplies are often configured for redundancy.

Flexibility

Open systems that use industry-standard architectures can be more flexible because users are not locked into a particular vendor. Another flexible option is remote support capability — both in-band (over a network) and out-of-band (via modem connection).

Useability

Hot-swap and server failover capabilities can allow upgrades and repairs with little or no downtime.

Accessibility

Thoughtful planning and engineering can have a large impact on serviceability. Using standardized hardware (nuts, bolts, screws) and making parts easy to reach can improve serviceability, thereby decreasing downtime.

Dedicated Servers

The main functional categories of servers are:

- File servers
- Print Servers
- Database server
- Application Servers
- Inter-Intranet Servers (specialized application servers)

Server Operating Systems

Key considerations when selecting a server operating system include:

- Compatibility with existing systems
- Ease of configuration and maintenance
- Scalability
- Number of concurrent users it will support

*UNIX** — UNIX was developed in 1969 by AT&T and was written in the "C" programming language. UNIX was ported to a wider variety of machines than any other operating system and is sold in different variations by several vendors.

Vendor	UNIX
Sun Microsystems*	Solaris*
SCO* (Santa Cruz Operations)	Sco UnixWare*
IBM*	AIX*
Digital*	OSF/1*
Hewlett-Packard*	HP/UX*

Windows NT Server — Windows NT is a 32-bit operating system introduced by Microsoft in 1993.

*NetWare** — NetWare is a 32-bit family of network operating systems from Novell. It supports DOS, Windows, OS/2* and Macintosh clients. UNIX client support is available from third parties.

Server Platforms

Some of the criteria for selecting the proper platform in a business and decision support environment include: complex processing power, the number of active users supported, amount of storage, availability, reliability, data integrity, security and disaster recovery.

Entry Level Servers — Entry Level Servers are relatively inexpensive and readily available from a variety of vendors. Entry level servers are the easiest to install and operate. They are ideal for small networks. They also are the least scalable.

Mid-range and Enterprise Servers — Mid-range and Enterprise servers often include fault tolerant features, scalability and multiple processors. Mid-range and Enterprise servers typically employ RAID subsystems and high-speed I/O subsystems for optimizing performance.

Which to Buy? — There are tradeoffs associated with implementing many small servers or fewer large servers. Consolidated servers pose a greater possibility of the entire network going down or business function being unavailable should the server or one of its components fail. However, larger servers are often loaded with more reliability, high availability and management features, so they may be less likely to fail. The maintenance on fewer larger servers should be less expensive than the maintenance on more smaller servers.

System software includes BIOS, operating systems and device drivers.

BIOS

BIOS stands for Basic Input Output System. It is the controlling layer between the computer operating system and the hardware. BIOS translates commands from software into signals that a computer needs to carry out the commands. It also ensures that the software works the same no matter which brand or model the BIOS is running on.

BIOS is located in Read Only Memory (ROM) components, usually found on the motherboard. It can be updated with a simple software upgrade.

Operating System (OS)

The operating system defines how applications work with the computer system. The OS establishes rules for applications to follow concerning files, memory allocation, and read/write operations.

The OS acts as a layer between the application software and the BIOS. The application makes a standard call to the OS, which in turn handles that task by communicating with the BIOS.

Device Drivers

Device drivers act like extensions of the operating system or BIOS by translating software commands into electrical signals needed to control hard drives, keyboards, printers, tape systems, memory and other peripherals. The operating system or BIOS calls the driver and the driver "drives" the device.

Device drivers are necessary because the OS and BIOS cannot be current enough to include all the code needed to control every existing hardware device. A device driver must be installed whenever any new hardware device is added to a system.

Part Number 283028-001

1098/HB/CS/PG